REMARKS

The Office Action dated December 29, 2004, has been received and carefully considered. In this response, claims 25-27 have been added. Entry of added claims 25-27 is respectfully requested. Reconsideration of the outstanding rejections in the present application is also respectfully requested based on the following remarks.

Applicants note with appreciation the indication on page 9 of the Office Action that claims 5, 6, 13, 14, 19 and 20 would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims. Applicants have opted to defer rewriting the above-identified claims in independent form pending reconsideration of the arguments presented below with respect to the rejected independent claims.

I. THE OBVIOUSNESS REJECTION OF CLAIMS 1-4, 7-12, 15-18, AND 21-24

On page 2 of the Office Action, claims 1, 2, 4, 5, 7-10, 12, 15-18 and 21-24 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Shinbashi *et al.* (U.S. Patent No. 5,796,717) in view of Albert *et al.* (U.S. Patent No. 6,606,315). On page 8 of the Office Action, claims 3 and 11 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Shinbashi *et*

al. in view of Albert *et al.* as applied to claims 9 and 1, and further in view of Adams, Jr. *et al.* (U.S. Patent No. 5,444,782). These rejections are hereby respectfully traversed.

Regarding independent claim 9, the Examiner asserts -- and Applicants agree -- that Shinbashi *et al.* fails to explicitly disclose "synchronizing means operatively connected to the primary node and the backup node for synchronizing the at least one backup node and the primary node."[1] The Examiner further asserts, however, that "Albert *et al.* explicitly disclosed such synchronizing means operatively connected to the primary node and the backup node for synchronizing the at least one backup node and the primary node. *(see Fig. 2A: Service Managers 241 and 242; and col. 1, line 66 - Col. 2, line 2)."* The Examiner also asserts that "[a]t the time of the invention, it would be obvious to a person of ordinary skill in the art to combine such synchronizing means operatively connected to the primary node and the backup node for synchronizing the at least one backup node and the primary node, as taught by Albert with Shinbashi, so that data can be immediately transferred throughout the stand-by unit without processing if a failure occurs at the

[1] On page 7 of the Office Action, the Examiner states that claim 1 is rejected for the same reason as claim 9 because the apparatus in claim 9 can be used to practice the method steps of claim 1. The Examiner also states that the subject matter of claim 17 is similar to that of claim 1, therefore the rejection of claim 1 would apply to reject the article of manufacture of claim 17 as well.

primary unit." The Examiner asserts the "motivation for doing so would have been to provide synchronization and control to eliminate the scalability limitations of the past in a data packet network *(se Albert: col. 9, lines 37-40)*. Therefore, it would have been obvious to combine Albert with Shinbashi

Applicants respectfully submit, however, that Albert does not teach or suggest "synchronizing means operatively connected to the primary node and the backup node for synchronizing the at least one backup node and the primary node." Rather, Applicants respectfully submit that Albert merely teaches two service managers 241 and 242 that provide decision-making capability (e.g., load balancing) to two forwarding agents:

> A service manager 241 and a second service manager 242 also communicate with the forwarding agents. ***The service managers provide the decision making capability that is required to provide a network service such as load balancing. The service managers send specific instructions to each of the forwarding agents detailing how certain flows of packets are to be processed. Such packet processing may include simply routing the packet, gathering statistics about the packet, sending the packet to a service manager, sending a notification that the packet has been seen to a service manager, modifying the packet, or using a special method such as tunneling or tag switching to send the packet to a destination other than the destination specified by the destination IP address included in the packet header.*** It should also be noted that forwarding agents in other embodiments also modify other aspects of packets, including packet source and destination addresses and port numbers and, in some instances, packet data.

See, Shinbashi, Col. 7, lines 22-39 (emphasis added).

Applicants respectfully submit that the above excerpt does not teach or suggest "synchronizing means operatively connected to the primary node and the backup node for synchronizing the at least one backup node and the primary node," as expressly recited in claim 9. Specifically, Applicants respectfully submit that there is no teaching or suggestion that service managers disclosed in Albert function to "synchronize" the two forwarding agents in the manner claimed.

Moreover, Applicant respectfully submits that one of ordinary skill in the art would not have been motivated to combine the teachings of Shinbashi and Albert. Shinbashi relates to a system for switching units in digital multiplexing equipment having a plurality of units for multiprocessing signals (e.g., multiplexing or de-multiplexing), and would thus not benefit from incorporating, and indeed is incompatible with, the load balancing and packet filtering systems and methods disclosed by Albert. In fact, even if the two references were combined, Applicants respectfully submit such a combination would not result in the claimed invention. In particular, Albert expressly teaches that "[t]here is no point through which all traffic between devices connected to network 210 and the group of servers must pass. ***Instead some traffic from network 210 and group of servers 220 passes through a forwarding agent***

231 and some traffic between network 210 and group of servers 220 passes through a forwarding agent 232." Applicant respectfully submits, therefore, that the emphasized portion makes clear that Albert's diffused distribution of signals functionality is not compatible with, and indeed works against, the back-up systems and methods disclosed by Shinbashi.

The remaining independent claims (e.g., claims 1 and 17) recite related subject matter to independent claim 9, and are therefore allowable for reasons similar to those given above.

The dependent claims 2-4, 7-8, 10-12, 15-16, 18, and 21-27, are allowable at least by virtue of their dependency on the above-identified independent claims. Moreover, these claims recite additional features which are not claimed, disclosed, or even suggested by the cited references taken either alone or in combination. For example, claims 25-27 expressly recite "wherein the ingress and egress traffic comprise session context information." Applicant respectfully submit that neither Shinbashi, Albert, nor Adam, alone or in combination, teach or suggest such a feature or functionality.

In view of the foregoing, it is respectfully requested that the aforementioned obviousness rejection of claims 1-4, 7-12, 15-18, and 21-24 be withdrawn.

II. CONCLUSION

In view of the foregoing, it is respectfully submitted that the present application is in condition for allowance, and an early indication of the same is courteously solicited. The Examiner is respectfully requested to contact the undersigned by telephone at the below listed telephone number, in order to expedite resolution of any issues and to expedite passage of the present application to issue, if any comments, questions, or suggestions arise in connection with the present application.

To the extent necessary, a petition for an extension of time under 37 CFR § 1.136 is hereby made.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-0206, and please credit any excess fees to the same deposit account.

Respectfully submitted,

Hunton & Williams LLP

By: 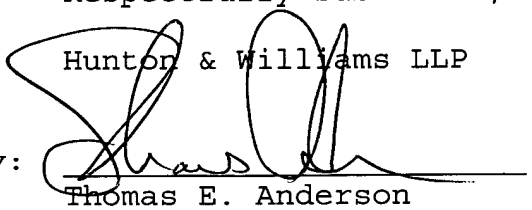

Thomas E. Anderson
Registration No. 37,063

TEA/OAF/dja

Hunton & Williams LLP
1900 K Street, N.W.
Washington, D.C. 20006-1109
Telephone: (202) 955-1500
Facsimile: (202) 778-2201
Date: April 29, 2005

APPENDIX A

1 (Previously Presented). A method for preventing information losses due to network node failure, the method comprising the steps of:

operatively connecting at least one backup node to a primary node;

synchronizing the at least one backup node and the primary node;

receiving, from a first endpoint, ingress traffic in the primary node;

replicating the ingress traffic to the at least one backup node;

outputting, from the primary node, primary egress traffic;

outputting, from the at least one backup node, backup egress traffic;

determining if the primary node has failed;

transmitting, to a second endpoint, the primary egress traffic if it is determined that the primary node has not failed; and

transmitting, to the second endpoint, the backup egress traffic from a selected one of the at least one backup nodes if it is determined that the primary node has failed,

wherein the backup egress traffic from the selected one of the at least one backup nodes replaces the primary egress traffic to the second endpoint and the backup node becomes the primary node for subsequent traffic.

2 (Original). The method of claim 1, wherein the primary node and the at least one backup node are network routers.

3 (Original). The method of claim 1, wherein the primary node and the at least one backup node are security engines for receiving encrypted ingress traffic and outputting decrypted egress traffic.

4 (Original). The method of claim 1, wherein the step of synchronizing the at least one backup node and the primary node further comprises the steps of:

transmitting synchronization information from the primary node to the at least one backup node.

5 (Original). The method of claim 4, wherein the step of transmitting synchronization information from the primary node to the at least one backup node further comprises the steps of:

transmitting at least one checkpoint message from the primary node to the at least one backup node, wherein the at least one checkpoint message includes static information

relating to the primary node as well as any outstanding session context for the primary node.

6 (Previously Presented). The method of claim 5, further comprising the steps of:

receiving, from the at least one backup node, a checkpoint message acknowledgment for each of said at least one checkpoint messages;

determining whether each of the checkpoint message acknowledgments was received prior to a change in flow state;

transmitting a synchronization declaration from the primary node to the at least one backup node if it is determined that each of the checkpoint message acknowledgments was received prior to a change in flow state; and

transmitting at least one new checkpoint message from the primary node to the backup node if it is determined that each of the checkpoint packet acknowledgments was not received prior to a change in flow state.

7 (Original). The method of claim 4, further comprising the steps of:

periodically assessing synchronization maintenance between the primary node and the at least one backup node.

8 (Original). The method of claim 7, wherein the step of periodically assessing synchronization maintenance further comprises the step of:

transmitting at least a portion of internal state information from the primary node to the at least one backup node sufficient to permit replication of primary node traffic on the at least one backup node.

9 (Original). An apparatus for preventing information losses due to network node failure, the apparatus comprising:

a primary node;

at least one backup node operatively connected to the primary node;

synchronizing means operatively connected to the primary node and the backup node for synchronizing the at least one backup node and the primary node;

means for receiving ingress traffic in the primary node from a first endpoint;

means for replicating the ingress traffic to the at least one backup node;

means for outputting primary egress traffic from the primary node;

means for outputting backup egress traffic from the at least one backup node;

determining means operatively connected to the primary node and the at least one backup node for determining whether the primary node has failed;

means for transmitting the primary egress traffic from the primary node to a second endpoint if the determining means determine that the primary node has not failed; and

means for transmitting the backup egress traffic from a selected one of the at least one backup nodes to the second endpoint if the determining means determine that the primary node has failed.

10 (Original). The apparatus of claim 9, wherein the primary node and the at least one backup node are network routers.

11 (Original). The apparatus of claim 9, wherein the primary node and the at least one backup node are security engines for receiving encrypted ingress traffic and outputting decrypted egress traffic.

12 (Original). The apparatus of claim 9, wherein the synchronizing means further comprise:

means for transmitting synchronization information from the primary node to the at least one backup node.

13 (Original). The apparatus of claim 12, wherein the means for transmitting synchronization information further comprise:

means for transmitting at least one checkpoint message from the primary node to the at least one backup node, wherein the at least one checkpoint message includes static information relating to the primary node as well as any outstanding session context for the primary node.

14 (Previously Presented). The apparatus of claim 13, further comprising:

means for receiving in the primary node, from the at least one backup node, a checkpoint message acknowledgment for each of said at least one checkpoint packets;

second determining means for determining whether each of the checkpoint message acknowledgments was received prior to a change in flow state;

means for transmitting a synchronization declaration from the primary node to the at least one backup node if it is determined that each of the checkpoint message acknowledgments was received prior to a change in flow state; and

means for transmitting at least one new checkpoint message from the primary node to the backup node if it is determined that each of the checkpoint message acknowledgments was not received prior to a change in flow state.

15 (Original). The apparatus of claim 12, further comprising:

means for periodically assessing synchronization maintenance between the primary node and the at least one backup node.

16. (Previously Presented) The apparatus of claim 15, wherein the means for periodically assessing synchronization maintenance further comprise:

means for transmitting at least a portion of an internal state of the primary node to the backup node sufficient to permit replication of primary node traffic on the at least one backup node.

17 (Original). An article of manufacture for preventing information losses due to network node failure, the article of manufacture comprising:

at least one processor readable carrier; and

instructions carried on the at least one carrier;

wherein the instructions are configured to be readable from the at least one carrier by at least one processor and thereby cause the at least one processor to operate so as to:

synchronize a primary node and at least one operatively connected backup node;

receive, from a first endpoint, ingress traffic;

replicate the ingress traffic to the at least one backup node;

output, from the primary node, primary egress traffic related to the ingress traffic;

output, from the at least one backup node, backup egress traffic related to the ingress traffic;

determine if the primary node has failed;

transmit, from the primary node, primary egress traffic related to the ingress traffic to a second endpoint if it is determined that the primary node has not failed; and

transmit, from a selected one of the at least one backup nodes, backup egress traffic to the second endpoint if it is determined that the primary node has failed,

wherein the backup egress traffic replaces the primary egress traffic to the second endpoint and the selected one of the at least one backup nodes becomes the primary node for subsequent traffic.

18 (Original). The article of manufacture of claim 17, wherein the instructions further cause the at least one processor to operate so as to:

transmit synchronization information from the primary node to the at least one backup node.

19 (Original). The article of manufacture of claim 18, wherein the instructions further cause the at least one processor to operate so as to:

transmit at least one checkpoint message from the primary node to the at least one backup node, wherein the at least one checkpoint message includes static information relating to the primary node as well as any outstanding session context for the primary node.

20 (Previously Presented). The article of manufacture of claim 19, wherein the instructions further cause the at least one processor to operate so as to:

receive, from the at least one backup node, a checkpoint message acknowledgment for each of said at least one checkpoint messages;

determine whether each of the checkpoint message acknowledgments was received prior to a change in flow state;

transmit a synchronization declaration from the primary node to the at least one backup node if it is determined that each of the checkpoint message acknowledgments was received prior to a change in flow state; and

transmit at least one new checkpoint message from the primary node to the backup node if it is determined that each of

the checkpoint message acknowledgments was not received prior to a change in flow state.

21 (Original). The article of manufacture of claim 18, wherein the instructions further cause the at least one processor to operate so as to:

periodically assess synchronization maintenance between the primary node and the at least one backup node.

22 (Original). A computer data signal embodied in a carrier wave readable by a computing system and encoding a computer program of instructions for executing a computer process performing the method recited in claim 1.

23 (Previously Presented). The method of claim 1 wherein the step of replicating the ingress traffic to the at least one backup node comprises simultaneously passing a copy of the ingress traffic to the at least one backup node.

24. (Previously Presented) The apparatus of claim 9 wherein the means for replicating the ingress traffic to the at least one backup node comprises means for simultaneously passing a copy of the ingress traffic to the at least one backup node.

25 (**New**). The method of claim 1 wherein the ingress and egress traffic comprise session context information.

26 (**New**). The apparatus of claim 9 wherein the ingress and egress traffic comprise session context information.

27 (**New**). The article of manufacture of claim 17 wherein the ingress and egress traffic comprise session context information.